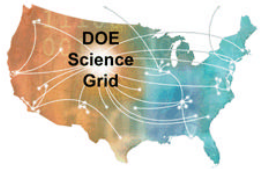

Scaling Issues for Large-Scale Grids: Security and Directory Services

William E. Johnston
Distributed Systems Department
NERSC Division

Michael Helm
DOE Energy Sciences Network

Lawrence Berkeley National Laboratory

Outline

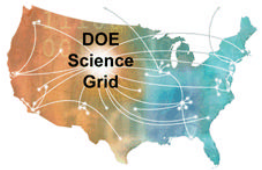

- Technology Issues for Large-Scale Grids
- Grid Directory Services
  - ESNet Can Play a Very Important Role in the Science Grid
- Security Aspects of Grids
  - ESNet Can Provide a Valuable Service for SciGrid Collaborations

➢ *Technology Issues for Large-Scale Grids*

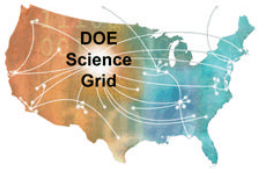

- Deployment of Grids to address large-scale science must accommodate the *virtual organizations* of the scientific collaborations
- Large virtual organizations have non-exclusive membership and must interact with other virtual organizations - at least in the realm of shared resources
- Two technologies needed to support this will be discussed here:
  - Grid Directory Services
  - Grid identity and security services

➢Grid Directory Services

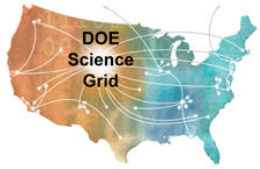

- The Grid will be a global infrastructure, and it will *depend heavily on the ability to locate information* about computing, data, and human resources for particular purposes, and within particular contexts.
- Most Grids will serve virtual organizations whose members are affiliated by some common characteristics
  - administrative parent (e.g. the DOE Science Grid and NASA's Information Power Grid)
  - long-lived project (e.g. a High Energy Physics experiment)
  - funding source, etc.
- Grid Directory Services refer to the implementation of the large-scale directory "hierarchies" needed for global scale Grids (as opposed to the GIS components and protocols that manage and make available information)

User Requirements

- Searching
  - The basic sort of question that a GDS/GIS must be able to answer is: for all resources in a virtual organization, provide a list of those with specified characteristics.
  - For example:
    - "Within the scope of the Atlas collaboration, return a list of all Sun systems with at least 2 CPUs and 1 gigabyte of memory, and that are running Solaris 2.6 or Solaris 2.7."
  - Answering this question involves determining the scope of the virtual organization ("Atlas"), querying all of the resources/lower layer directory servers, and then "filtering" resource attributes in order to produce a list of candidates.

Grid Directory Services: User Requirements - Virtual Organizations -

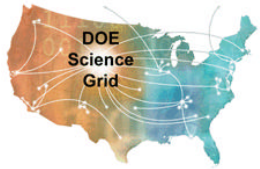

- The Grid Directory Service (GDS) should be able to provide “roots” for virtual organizations.
  - These nodes provide search scoping by establishing the top of a hierarchy/index node of virtual org. resources, and therefore a starting place for searches.
  - Like other named objects in the Grid, these virtual org. nodes might have characteristics specified by attributes and values.
    In particular, the virtual organization node probably needs a name reflecting the org. name, however some names (e.g. for resources) may be inherited from the Internet DNS domain names.

Grid Directory Services

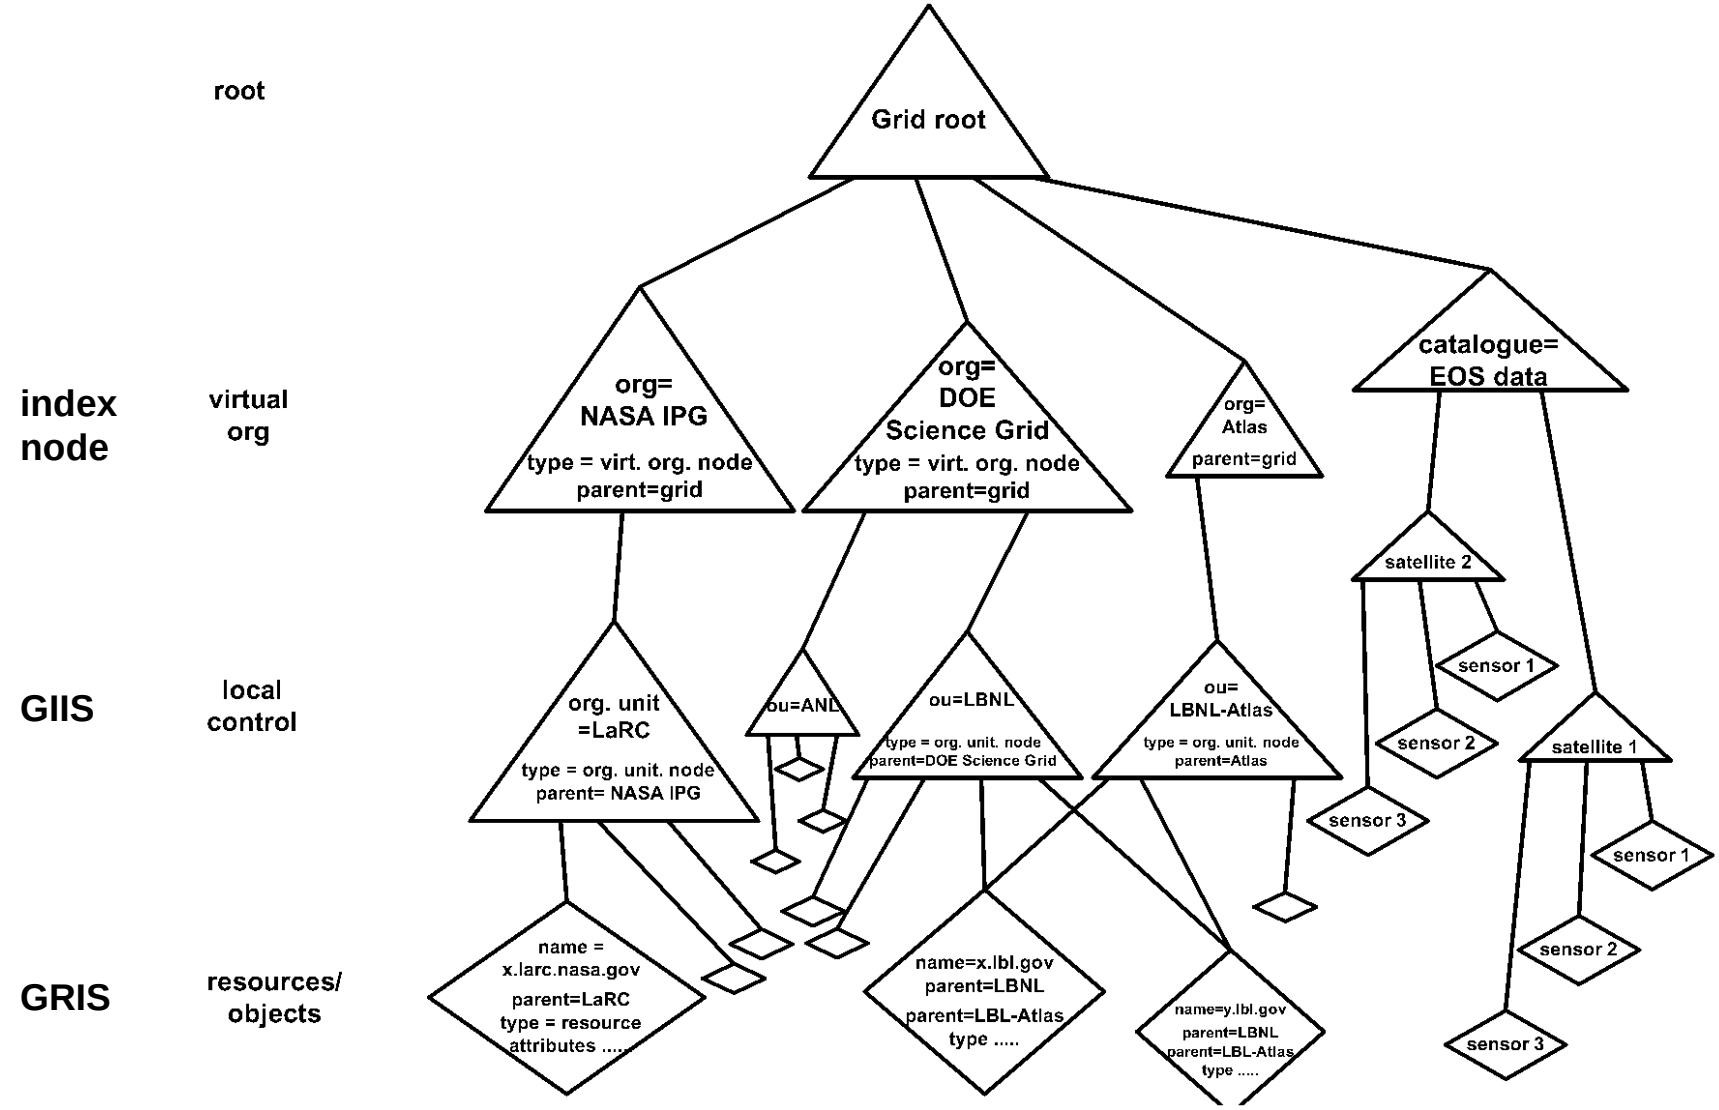

Grid Directory Services: User Requirements - Information and Data Objects -

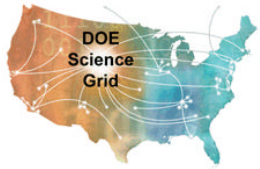

- A variety of other information will require cataloguing and global access, and the GDS should accommodate this in order to *minimize the number of long-lived servers that have to be managed*. E.g.:
  - metadata catalogues
  - dataset replica information
  - database registries
  - Grid monitoring objects (system state, user allocations, nets, etc.)
  - Grid entity certification/registration authorities (e.g. X.509 Certificate Authorities)
  - Grid Information Services object schema

Operational Requirements

- Performance, Reliability, Operability
  - Queries, especially local queries, should be satisfied in times that are comparable to other local queries, such as uncached DNS data for local systems. E.g., seconds or fractions of seconds – this is a substantial scaling issue
  - Local sites should not be dependent on remote servers to locate and search local resources.
  - It should be possible to restrict searches to resources of a single, local, administrative domain.
  - Minimal manual management - e.g., lower-level objects should auto register with directory servers - a “self organizing” structure

Operational Requirements

- Site administrative domains may wish to restrict external access to local information, and therefore will want control over a local, or set of local, information servers.
  - This implies the need for servers intermediate between local resources and the virtual org. root node that are under local control for security, performance management, and reliability management.
  - (Note that in the Globus terminology that these intermediate directory servers are called GIISs.)

Grid Directory Services:

Operational Requirements

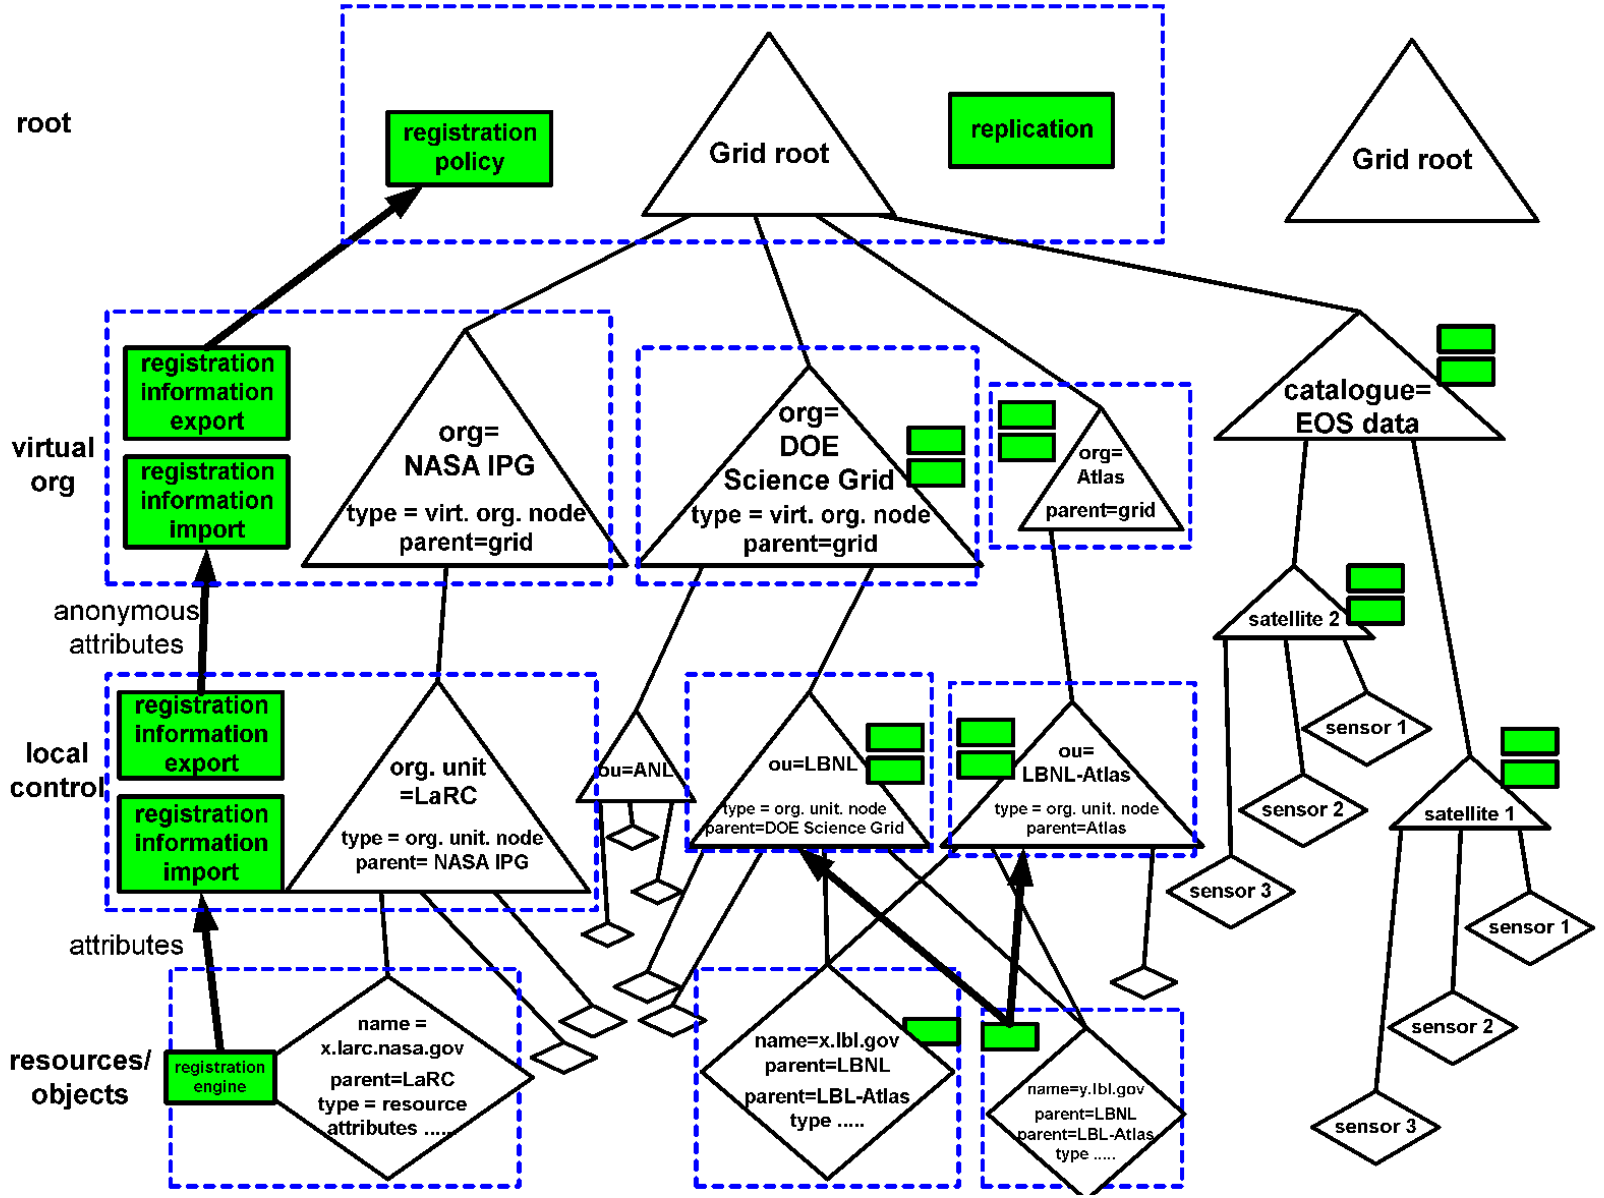

The Importance of Grid Directory Services for a Scalable Science Grid

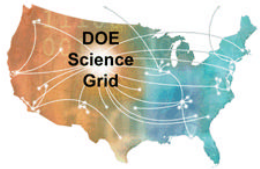

- Provide a coherent view of a large infrastructure
  - must be a higher level namespace manager, or you get an N X N problem with no guarantee of naming consistency
  - GIS “index nodes” are directory servers that provide customized views lower-level domains
    - this is fine for a group of related projects, but every such group must set up their own resulting in an ad hoc and probably incomplete view of the Grid
    - unless DNS names are used (and this is not always suitable) there is no guarantee of a consistent namespace

Grid Directory Services: Ad Hoc Webs and Managed Hierarchy

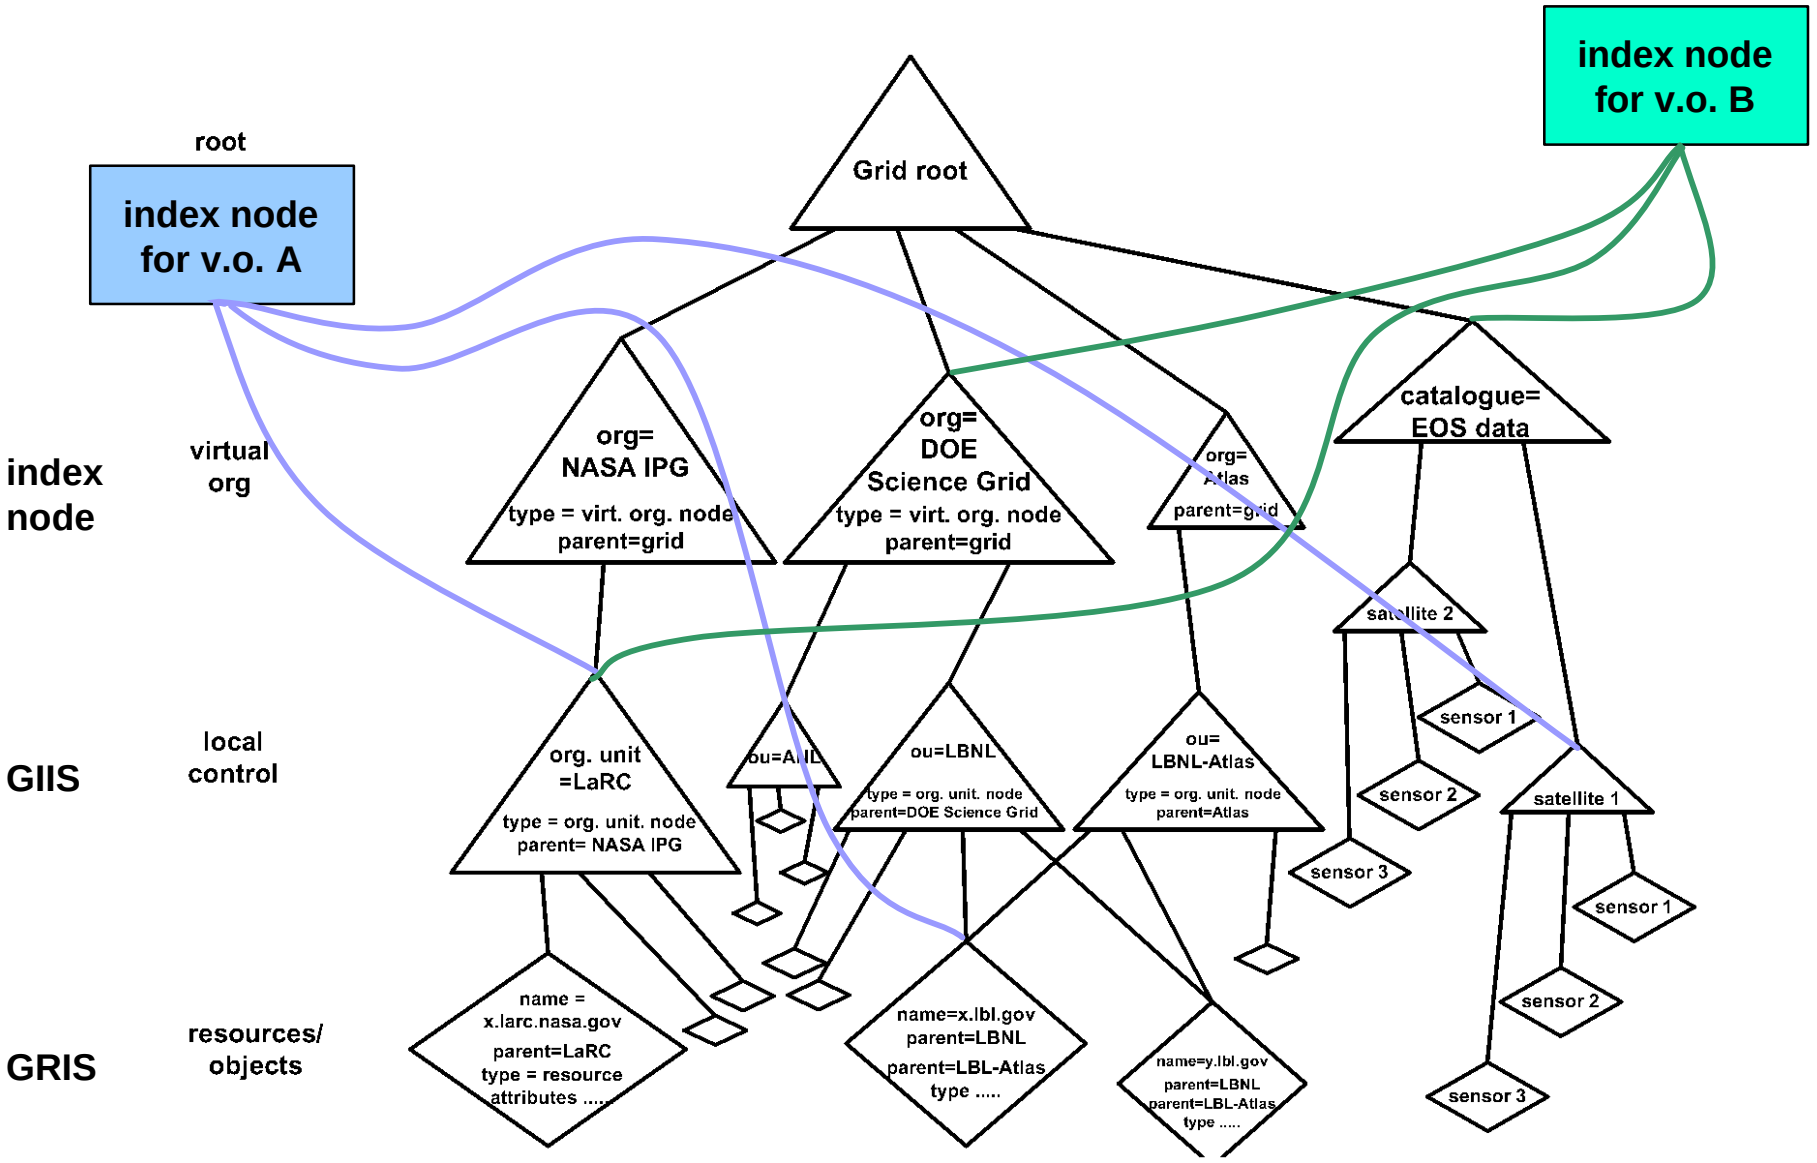

Grid Directory Services

- Both hierarchical and index node/web directory services are useful, but in the Science Grid context, a managed hierarchy will be important and very useful for managing large-scale virtual org. structures

•ESNet Can Play a Very Important Role in the Science Grid

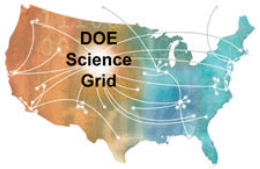

- ESNet can provide a rooted and managed namespace, and a place to home SciGrid and SciDAC virtual organizations
  - ESNet has the experience and credibility to do this - an could thus contribute substantially to the utility of the SciGrid
- ESNet could provide a valuable service to SciGrid based collaborations
  - Having to set up and operate a directory server is a large overhead for even sizable virtual organizations - metadirectory servers, operational requirements, long-lived servers, etc.

ESNet Can Play a Very Important Role in the Science Grid

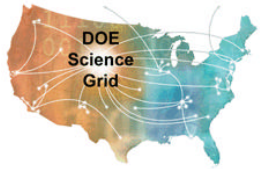

- In the longer term, ESNet can address the directory services scaling issues for large virtual orgs. - it will be in a unique position (by way of operating the higher-level directory servers) to see the global issues
- The Grid is inherently network based - it is arguably a network service - so all DOE Grid participants are also ESNet customers

➢Security Aspects of Grids

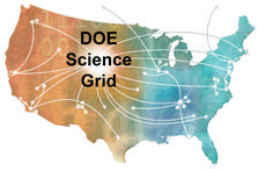

- Users are no longer listed in a single central database at a local site, however positive identification to an entity that can provide human accountability will still be required
  - Strong authentication to a globally unique identity via cryptographic credentials
  - Grid cryptographic credentials should have a well understood (published & audited) relationship to the human subject
  - Certification Authorities that have appropriate operating policy and issue X.509 identity certificates provide this service

Security Aspects of Grids

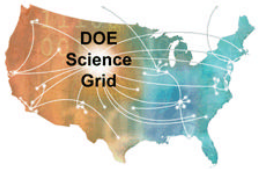

- Most Grid resources are remote from the users, so delegation of authority is needed - i.e. the user's authenticated identity must be carried to remote resources without the real-time intervention of the human user
  - this is accomplished with restricted delegation - "proxy" - credentials that are presented by the Grid Security Infrastructure for authentication and authorization at a remote resource

Security Aspects of Grids

- There will be multiple stakeholders for the resources involved in Grid applications who will probably not have a uniform resource use policy: Users may have to be authorized separately for every resource that in incorporated into a Grid application system.
  - Strong, flexible, and easily used and managed policy based authorization and access control (an R&D topic - see Akenti and GAA)

Security Aspects of Grids

- Grid services should not weaken security of local systems, and a security compromise on one platform that is involved in a Grid application system should not propagate via Grid services to other platforms in the system.
  - careful management of identity credentials
  - understanding what credential management clients are trusted (and untrusted!)
  - require the use of authenticated control channels (GSI services) between distributed application components

Security Aspects of Grids

- Grid users will not have control over the security policy of remote resources (e.g. computing platforms)
  - It may be necessary to “rate” systems on their security, and provide that rating as a system characteristic that may be used in choosing resources from a candidate pool when constructing the resource base for a distributed application.

Security Aspects of Grids

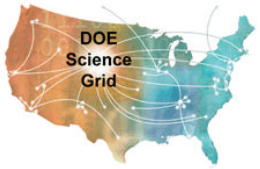

- The GDS and GIS must have security model of their own to address the confidentiality and integrity of the details of site resources, while at the same time permitting the sorts information searches needed to assemble application systems from many scattered resources.

The ***Certification Authority*** plays a central role in providing credentials needed for a system/resource to "trust" a remote user entity

- Establishes the "nature" of the binding of human identity to cryptographic token (the X.509 identity certificate binds a distinguished name to the public key part of a public/private key pair)
- Provides for recovery from loss of control of the identity token
- Operates with an understood level of integrity for the CA functions
- A "root CA" plays an important role in the integrity of the organizational CAs

Identity Certification Authority

- "Nature" of the binding of human identity to cryptographic token (the X.509 identity certificate)
  - This is governed by a Certificate Policy that specifies who the CA will certify and under what conditions
  - "who" is typically members of an organization or virtual org.
  - "under what conditions" specifies how human identity must be established before a certificate is issued

Security Aspects of Grids:
Identity Certification Authority

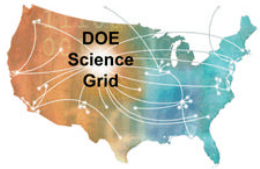

- Provides for recovery from loss of control of the identity token
  - must maintain and publish accurate records about the state of the certificates that it issues: valid, expired, suspended, revoked
  - when a user looses control of an identity certificate – compromise of the passphrase - it must be revoked and a new one issued
    - the use of an identity certificate is controlled by a passphrase that is used to decrypt the private key so that it may be used to encrypt something to prove that the certified user is the originator -
  - revoking and reissuing a certificate is “equivalent” to ensuring that a user’s password has been changed
  - typically done by publishing a Certificate Revocation List that may be checked when a certificate is presented as proof of identity

Identity Certification Authority

- Operates with an understood level of integrity for the CA functions
  - the CA produces an identity certificate by combining the user's Distinguished Name and public key in a standard format (X.509) and then digitally signing this document
  - all aspects of the CA signing function must be carefully protected
    - off-line
    - small number of people have the CA's passphrase
    - bogus certificates must not be signed
    - all actions must be logged for auditing
    - these practices are codified in the Certificate Policy and verified by some form of audit

Identity Certification Authority

- A “root CA” plays an important role in the integrity of the organizational CAs
  - The CA itself has an identity certificate that is used to verify the certificates that it issues
  - The CA certificate is typically self-signed - the CA signs it using its own private key
    - if the CA private key is compromised, then a bogus CA certificate could be generated - if this cert is signed by a higher level (e.g. root) CA, then this is much more difficult
  - A root CA might also represent the validity of CP audits
    - an organizational CA’s certificate could be revoked (placed on the root CA’s CRL) for failing to provide audits
  - Root CA should provide a repository of CA certs. and CRLs

Virtual Organizations and Certification Authorities

- Both real and virtual organizations are likely to have differing requirements for Certificate Policies
- Even if the CPs are functionally equivalent, there may be requirements that differ in detail
- There may be significant differences in CP, in which case a hierarchy of “strictness” would be useful in evaluating trust

Virtual Organizations and Certification Authorities

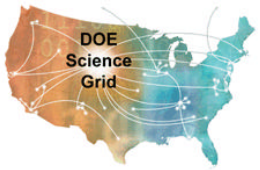

- The Science Grid will serve many different virtual organizations with different CP requirements, yet a centrally managed CA will be a valuable service (correct operation of a CA within a CP is hard)

•ESNet Can Provide a Valuable Service for SciGrid Collaborations

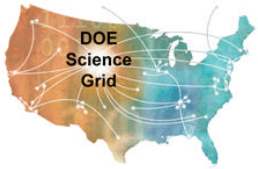

- As with directory services, ESNet is the logical provider of CA functionality, which is critically important for the Science Grid and other SciDAC projects that use PKI based security
  - ESNet has the experience and credibility to do this - an could thus contribute substantially to the utility of the SciGrid
  - ESNet has a broad customer base within DOE
  - ESNet has a history of success with similar services
- Having to set up and operate a CA is a large overhead for even sizable virtual organizations - CPs, operational requirements, long-lived servers, etc.

ESNet and the DOE Science Grid CA: A Proposal

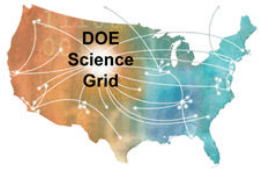

- ESNet will run the SciGrid root CA, which will sign CA certs for any CP compliant, SciDAC and/or SciGrid related virtual organization
- ESNet will run the SciGrid signing CA
  - this will be several logical CAs that represent different virtual organizations and therefore different CPs
  - it will operate at a sufficiently rigorous level to accommodate all of the operational requirements of the v.o. CPs
  - for each v.o. / CP there will be a registration agent that is responsible for carrying out the identity verification requirements of that v.o.'s CP
  - will maintain a CRL for each v.o. logical CA

ESNet and the DOE Science Grid CA: A Proposal

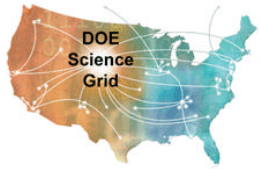

- ESNet will run the CA certificate and CRL repositories
- All of this should fall largely within current, standard ESNet operational practices for its infrastructure (physical security, established operational procedures, careful system management and administration)
- All of the contact with the end users would be handled by the registration agents

ESNet and the DOE Science Grid CA: A Proposal

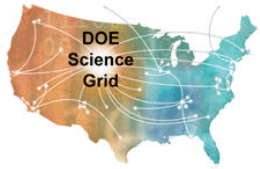

- In the longer term, the DOE Public Key Infrastructure may be able to provide part of this service in the long term, but not in the short term
  - technical issues: Entrust integration with Grid security
  - policy issues: a policy designed for Federal employees will not easily be extended to cover all SciDAC participants (in fact probably cannot be)

Acknowledgements

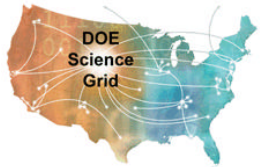

- This work is supported by the U.S. Dept. of Energy, Office of Science, Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Sciences Division under contract DE-AC03-76SF00098 with the University of California.

References and Acronyms

- www.globus.org provides full information about the Globus system.
- The Grid: Blueprint for a New Computing Infrastructure, edited by Ian Foster and Carl Kesselman. Morgan Kaufmann, Pub. August 1998. ISBN 1-55860-475-8. http://www.mkp.com/books_catalog/1-55860-475-8.asp
- “Grids as Production Computing Environments: The Engineering Aspects of NASA's Information Power Grid,” William E. Johnston, Dennis Gannon, and Bill Nitzberg. Eighth IEEE International Symposium on High Performance Distributed Computing, Aug. 3-6, 1999, Redondo Beach, California. (Available at http://www.ipg.nasa.gov “About IPG” -> presentations)
- “Vision and Strategy for a DOE Science Grid” - http://www.itg.lbl.gov/~wej/Grids
- See www.ipg.nasa.gov for project information and pointers.
- See http://www-itg.lbl.gov/NGI/ for project information and pointers.
- The Particle Physics Data Grid has two long-term objectives. Firstly: the delivery of an infrastructure for very widely distributed analysis of particle physics data at multi-petabyte scales by hundreds to thousands of physicists. Secondly: the acceleration of the development of network and middleware infrastructure aimed broadly at data-intensive collaborative science. http://www.cacr.caltech.edu/ppdg/
- “The Data Grid: Towards an Architecture for the Distributed Management and Analysis of Large Scientific Datasets.” A. Chervenak, I. Foster, C. Kesselman, C. Salisbury, S. Tuecke, (to be published in the Journal of Network and Computer Applications).

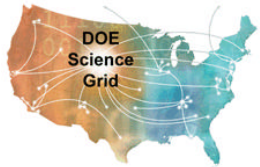

- The Grid Forum (www.gridforum.org) is an informal consortium of institutions and individuals working on wide area computing and computational Grids. Current working groups include Security (authentication, authorization), Scheduling and Resource Management, Grid Information Services, Application and Tool Requirements, Advanced Programming Models, Grid User Services and Operations, Account Management, Remote Data Access, Grid Performance
- Akenti: “Certificate-based Access Control for Widely Distributed Resources,” Mary Thompson, William Johnston, Srilekha Mudumbai, Gary Hoo, Keith Jackson, Usenix Security Symposium ‘99. Mar. 16, 1999. (See http://www-itg.lbl.gov/Akenti)
- GAA: “Generic Authorization and Access control API” (GAA API). IETF Draft. http://ghost.isi.edu/info/gss_api.html)
- Storage Resource Broker (SRB) provides uniform access mechanism to diverse and distributed data sources. http://www.sdsc.edu/MDAS/
- Condor is a High Throughput Computing environment that can manage very large collections of distributively owned workstations. http://www.cs.wisc.edu/condor/
- SCIRun is a scientific programming environment that allows the interactive construction, debugging and steering of large-scale scientific computations. http://www.cs.utah.edu/~sci/software/
- Ecce - www.emsl.pnl.gov
- WebFlow - A prototype visual graph based dataflow environment, WebFlow, uses the mesh of Java Web Servers as a control and coordination middleware, WebVM. See http://iwt.npac.syr.edu/projects/webflow/index.htm
- “A CORBA-based Development Environment for Wrapping and Coupling Legacy Codes,” Gregory Follen, Chan Kim, Isaac Lopez, Janche Sang and Scott Townsend. To be presented at Tenth IEEE International Symposium on High Performance Distributed Computing, August 7-9, 2001, San Francisco, California.

- "NetLogger: A Toolkit for Distributed System Performance Analysis," D. Gunter, B. Tierney, B. Crowley, M. Holding and J. Lee. In *IEEE Mascots 2000: Eighth International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*. 2000. http://www-didc.lbl.gov/papers/NetLogger.Mascots.paper.ieee.pdf
- "Pipechar Network Characterization Service," G. Jin.
  Tools based on hop-by-hop network analysis are increasingly critical to network troubleshooting on the rapidly growing Internet. Network characterization service (NCS) provides ability to diagnose and troubleshoot networks hop-by-hop in an easy and timely fashion. http://www-didc.lbl.gov/NCS/